

IN THE CLAIMS:

The current claims follow. For claims not marked as amended in this response, any difference in the claims below and the previous state of the claims is unintentional and in the nature of a typographical error.

1. (Previously Presented) An address lookup structure comprising:
a plurality of hash tables each storing prefixes for address lookups;
a content addressable memory storing at least some prefixes for which a collision occurs within at least one of the hash tables; and
a hashing lookup search mechanism that comprises:
a routing table implemented with selective hashing for a plurality of prefixes with different lengths; and
a plurality of memory blocks, wherein each hash table is allocated a group of the memory blocks based on a size of the respective hash table.

2. (Previously Presented) The address lookup structure according to claim 1, wherein each hash table is allocated a smallest number of memory blocks sufficient to hold prefixes for which no collision occurs within the respective hash table.

3. (Previously Presented) The address lookup structure according to claim 1, wherein each hash table is allocated no more than a predefined number of memory blocks.

4. (Previously Presented) The address lookup structure according to claim 1, wherein each hash table contains prefixes hashed by one of two hash functions, a second of the two hash functions employed when a collision occurs with a first of the two hash functions.

5. (Previously Presented) The address lookup structure according to claim 1, wherein each hash table contains different length prefixes.

6. (Previously Presented) The address lookup structure according to claim 5, further comprising:

a priority encoder selecting a longest prefix when a plurality of matches occur between different length portions of a prefix and prefixes in each of two or more of the hash tables.

7. (Original) The address lookup structure according to claim 5, wherein the plurality of hash tables contain only a subset of different length prefixes possible under an addressing scheme, and wherein a remainder of the different length prefixes are stored in the content addressable memory.

8. (Previously Presented) A network router including the address lookup structure according to claim 1, the network router further comprising:

a network search engine containing the hash tables and coupled to the content addressable memory, the network search engine performing address lookups using the hash tables; and

an external memory coupled to the network search engine and containing per route information indexed by a next hop index generated by the network search engine.

9. (Original) A network including a plurality of interconnected network routers according to claim 8.

10. (Previously Presented) An address lookup structure comprising:
a plurality of hash tables each containing prefixes of a different length than prefixes within other hash tables within the plurality, the hash tables collectively containing only a subset of different prefix lengths less than or equal to an address length;
an additional address lookup facility handling a remainder of the different address lengths not accommodated by the plurality of hash tables; and
a hashing lookup search mechanism that comprises:
a routing table implemented with selective hashing for a plurality of prefixes with different lengths; and
a plurality of memory blocks, wherein each hash table is allocated a group of the memory blocks based on a size of the respective hash table.

11. (Original) The address lookup structure according to claim 10, wherein the additional address lookup facility comprises a content addressable memory.

12. (Previously Presented) The address lookup structure according to claim 10, wherein each of the plurality of hash tables is contained in one or more memory blocks allocated based on hashing of each prefix contained in the respective hash table using at least a first hash function,

wherein a number of memory blocks allocated to the respective hash table does not exceed a predefined number, and

wherein a remainder of prefixes of a length corresponding to prefixes within the respective hash table are handled by the additional address lookup facility.

13. (Original) The address lookup structure according to claim 10, further comprising:

a priority encoder selecting a longest prefix match from matches identified within the plurality of hash tables.

14. (Previously Presented) A method of operating an address lookup comprising:
storing at least some address prefixes in each of a plurality of hash tables, wherein each hash table is allocated a different group of memory blocks from a plurality of memory blocks, wherein a number of the memory blocks allocated to a hash table is based on a size of the respective hash table;

storing address prefixes for which a collision occurs within at least one of the hash tables in a content addressable memory; and

operating a hashing lookup search mechanism that comprises a routing table implemented with selective hashing for a plurality of prefixes with different lengths.

15. (Previously Presented) The method according to claim 14, further comprising:

maintaining each hash table within a smallest number of memory blocks sufficient to hold all required prefixes for which no collision occurs within the respective hash table.

16. (Previously Presented) The method according to claim 14, further comprising:

allocating each hash table no more than a predefined number of memory blocks.

17. (Previously Presented) The method according to claim 14, further comprising:

hashing prefixes in each hash table with one of two hash functions, a second of the two hash functions employed when a collision occurs with a first of two hash functions.

18. (Original) The method according to claim 14, further comprising:
storing, in each of a plurality of hash tables, prefixes of a different length than prefixes contained in any other of the plurality of hash tables.

19. (Original) The method according to claim 18, further comprising:
selecting a longest prefix when a plurality of matches occur between different length portions of a prefix and prefixes in each of two or more of the plurality of hash tables.

20. (Previously Presented) The method according to claim 18, further comprising:
storing prefixes corresponding to only a subset of different prefix lengths possible under an addressing scheme in the plurality of hash tables; and
storing a remainder of prefixes in the content addressable memory.

21. (Currently Amended) The address lookup structure according to claim 1, wherein the hashing lookup search mechanism further comprises:
a configuration register associated with each memory block, each configuration register identifying a prefix length to which the respective memory block is allocated.

22. (Currently Amended) The address lookup structure according to claim 1, wherein each configuration register identifies the hash function to which the respective memory block is allocated.